THE JOURNAL OF FINTECH & DIGITAL ASSETS

JFDA, Volume 1 (2), 2026

| Research Article

# Two-Regime Liquidity Recovery After a Perpetual Futures Liquidation Cascade: Evidence from Hyper-liquid and the October-10th-2025 Event.

**Lim Boon Chuan***

**Authors**
Lim Boon Chuan, Independent Researcher, Singapore.
Email: booncuan@singapore.to

* Corresponding author.

**Abstract**
We document the inside-market response of Hyperliquid - the largest decentralized perpetual futures venue - to the October 10, 2025 system-wide liquidation cascade. Using the venue's public L2 order-book archive, we construct a minute-level panel covering BTC, ETH, SOL, XRP, and DOGE perpetual contracts across a 92-day window. There are three main findings. First, the cascade was a tightly compressed intraday event. Median quoted spread pooled across assets reached 9.4 basis points at hour 21 UTC, and the 95th percentile reached 87.5 bps, with intraday recovery well under way by end of day. Second, across the 23-day post-event window prior to a separate market-wide selloff, quoted spreads on all five contracts returned to within 1.2 times pre-event levels, and three of five to within 1.1 times. Third, inside-quote depth contracted persistently over the same window, with four of five contracts ending between 66 and 79 percent of pre-event levels. We interpret these results as two-regime recovery, in which pricing recovers while inventory commitment persistently contracts. Spread-only metrics systematically understate post-event impairment of venue liquidity.

**Keywords:** Decentralized exchange, Market microstructure, Perpetual futures, Liquidation cascades, Market resilience
**JEL Classification:** G14, G15, G23

**Article Information**
Received: 24 April 2026
Revised: 16 June 2026
Accepted: 20 June 2026
Published: September 2026

**Citation:**
Lim, B. C. (2026). Two-regime liquidity recovery after a perpetual futures liquidation cascade: Evidence from Hyperliquid and the October-10th-2025 event. *The Journal of FinTech and Digital Assets*, 1(2), 54–68.

## 1. Introduction

Cryptocurrency derivatives have grown from a marginal segment into the largest and most actively traded part of the digital-asset market. Trading in these instruments now routinely exceeds spot trading in notional volume, and the segment spans futures, options, and perpetual contracts listed across both centralized and decentralized venues. Among these instrument types, the perpetual future - a futures-style contract with no expiry, held in line with the underlying through a periodic funding payment - has become the dominant vehicle for leveraged exposure to cryptocurrencies. Against this backdrop, we focus on the perpetual-futures segment. Perpetual futures are the dominant derivative class in cryptocurrency markets, with global notional volumes that have exceeded sixty trillion US dollars annually. Until recently, all meaningful liquidity in this segment resided on centralized exchanges. Hyperliquid, launched in 2023 and scaled materially through 2024 and 2025, is the first decentralized venue to attain comparable order-book liquidity. Its market share of decentralized perpetual volume reached approximately 60-65 percent by mid-2025, and its single-venue daily volume regularly exceeds five billion US dollars. The protocol's distinguishing feature is a consensus-layer ordering rule. Within each block, the validator software is required to process cancel orders and post-only limit orders before processing good-til-cancel and immediate-or-cancel orders. The rationale is grounded in microstructure theory. When prices move, market makers want to cancel stale quotes, and if a faster taker is processed first the maker is picked off. Cancel prioritization eliminates this specific source of adverse selection on resting orders by giving the maker a structural ordering advantage.

On October 10, 2025, the cryptocurrency derivatives market experienced a system-wide liquidation cascade. Across centralized and decentralized perpetual venues, leveraged positions were liquidated at unusual scale; aggregate liquidation volume exceeded nineteen billion US dollars within a 24-hour window. The cascade has been the subject of subsequent industry commentary and a formal analysis of the protocol's automated deleveraging mechanism (Chitra, 2025). What has not been documented is the inside-market response on Hyperliquid in detail - how quickly spreads widened, how quickly they recovered, and whether the recovery was symmetric between quoted prices and posted sizes.

This paper provides that documentation. Using Hyperliquid's public L2 order-book archive, we construct a minute-level panel of inside-market metrics for the five most-traded perpetual contracts (BTC, ETH, SOL, XRP, and DOGE) across a 92-day window centered on the October 10 event. The dataset comprises 531,305-minute observations. Because the entire analysis is venue-internal, all comparisons are against the venue's own pre-event baseline; no cross-venue spread inference is required.

There are three main findings obtained from the analysis. First, the cascade was a tightly compressed intraday event. Hour-by-hour analysis of October 10 UTC shows the inside spread remaining at normal levels (0.37 to 0.44 basis points pooled across all five assets) through hour 14, beginning to widen at hour 15, and peaking at hour 21 (median 9.38 bps, p95 87.48 bps, maximum 241.8 bps). By hour 23, median spread had declined to 1.08 bps. The bulk of the venue's stress event was contained within a three-hour window.

Second, spread recovery across the 23-day window between the cascade and a separate market-wide selloff beginning November 3, 2025 was close to complete. By November 2, 2025, daily median spreads returned to between 1.06- and 1.19-times pre-event baseline, with three of five contracts within 1.10 times. A subsequent market-wide selloff pushed spread trajectories back into the 1.3 to 1.5 range; we treat this second event explicitly in Section 4 and demonstrate that the clean-window recovery finding is robust to its exclusion.

Third, inside-quote depth contracted persistently over the same clean window, even as spreads recovered. Four of five contracts showed further depth contraction from the first week to the third week post-event, with end-of-window depth ratios between 0.66 and 0.79 of pre-event levels. The depth contraction preceded the November selloff and is therefore attributable to the cascade itself.

We term this pattern two-regime recovery, in which pricing recovers quickly and nearly completely while inventory commitment contracts persistently. The pattern is consistent with the inventory-management framework (Ho & Stoll, 1981; Stoll, 1978), in which a tail event raises perceived inventory risk and the rational maker response is size reduction rather than price markup. For execution-quality measurement, the implication is that spread-only metrics systematically understate post-event impairment of venue liquidity.

These findings contribute to two literatures. First, to the empirical microstructure of decentralized derivatives venues, which until recently has been data constrained. The public archive that Hyperliquid publishes makes inside-market analysis at this granularity feasible from outside the venue, at minimal cost. Second, to the literature on market resilience after stress events, where almost all empirical work has focused on equity venues (Anand & Venkataraman, 2016; Bessembinder et al., 2015; Brogaard et al., 2018). We provide what we believe is the first detailed inside-market resilience study on a major perpetual futures venue, centralized or decentralized.

The remainder of the paper proceeds as follows. Section 2 reviews related literature. Section 3 describes the institutional setting, data, and methodology. Section 4 presents empirical results and discussion, including the cascade event study, recovery dynamics, depth contraction findings, and robustness analyses. Section 5 concludes.

## 2. Literature Review

### 2.1. Cryptocurrency Market Microstructure

The microstructure of cryptocurrency markets has been an active research area since the emergence of derivatives trading on these assets. Foley et al. (2019) document order-flow dynamics on early Bitcoin spot exchanges. Brauneis et al. (2021) provide a comprehensive treatment of liquidity measurement in cryptocurrency markets, comparing transaction-cost proxies against high-frequency benchmarks across major exchanges. He et al. (2024) study perpetual-spot arbitrage and the associated execution costs. Ruan and Streltsov (2024) analyze how the introduction and removal of perpetual contracts affect spot-market liquidity and price efficiency. None of these papers analyze inside-market behavior on a decentralized perpetual venue, which is the focus of our work.

On-chain venues introduce distinct microstructure features absent from centralized exchanges, namely deterministic block-level execution, validator co-location, and protocol-defined order priorities. Chitra (2025) provide the first formal model of automated deleveraging mechanisms for perpetual exchanges and analyze them empirically using the Hyperliquid October 10 cascade. Their work establishes a trilemma between solvency, revenue, and trader fairness in the design of forced-liquidation mechanisms. The present paper is complementary, analyzing the inside-market response to the same event from the perspective of voluntary market-maker behavior rather than protocol-enforced liquidations.

### 2.2. Market Resilience After Stress Events

A substantial literature studies how equity markets respond to and recover from stress events. Brogaard et al. (2018) document liquidity dynamics around extreme price moves on US equity venues, finding that high-frequency traders supply liquidity during ordinary stress but withdraw during extreme stress. Anand and Venkataraman (2016) study recovery dynamics after market-wide circuit-breaker triggers. Bessembinder et al. (2015) examine the persistence of spread widening following corporate-event-driven volatility. The shared finding across these papers is that recovery is asymmetric. Spreads widen quickly in stress and narrow more slowly, with the slowness driven by maker risk-aversion rather than by underlying volatility persistence.

A common limitation of this literature is that it focuses almost exclusively on the spread dimension of execution quality, with depth treated as a secondary consideration or omitted entirely. The two-regime finding documented here suggests this focus misses the more persistent dimension of post-event impairment in at least one class of venue. To our knowledge no paper in this literature has examined a perpetual futures venue, and no paper has examined a decentralized venue. The Hyperliquid October 10 cascade is well-suited to extending this literature because it occurred on a venue with continuous (24/7) trading and no circuit breakers, providing a clean observation of maker-driven recovery dynamics absent venue-imposed interventions.

### 2.3. Inventory Management and Maker Responses to Risk

The theoretical foundation for the two-regime interpretation is the inventory-management framework of Stoll (1978) and Ho and Stoll (1981), in which market makers set quotes as a function of both expected adverse selection and the cost of holding inventory exposed to price risk. In this framework, a tail event that raises the perceived variance of inventory can rationally produce a separable response, in which competitive quotes are maintained to capture order flow while the size at those quotes is reduced to limit inventory accumulation. The two-regime pattern documented here is consistent with this model. The literature has previously documented the model's predictions in other settings (Hasbrouck & Sofianos, 1993; Madhavan & Smidt, 1991), but we are not aware of a direct decomposition of post-event recovery into pricing and inventory dimensions separately for a major perpetual futures venue.

### 2.4. Cancel Prioritization and Maker Protection

The literature on order-priority rules is long but has been relatively quiet on cancel prioritization specifically. Hendershott and Moulton (2011) document that the introduction of NYSE's Hybrid Market — which reduced the time-priority advantage of slow-moving liquidity providers — measurably worsened execution quality for retail-sized orders, suggesting that protections for resting orders matter for inside-market quality. Foucault et al. (2013) provide a theoretical treatment of cancel-related fees and their effect on quote-update behavior. Hyperliquid's design implements a particularly aggressive form of cancel prioritization, under which all cancels are processed at the consensus layer before any matching, regardless of submission time. To our knowledge no centralized exchange operates this rule in its full form.

## 3. Data and Methodology

### 3.1. Institutional Setting: Hyperliquid

Hyperliquid is a layer-one blockchain whose primary application is a perpetual-futures decentralized exchange. The protocol uses HyperBFT, a modified HotStuff consensus algorithm with sub-second block times. All order-book operations - placement, modification, cancellation, and matching - execute on-chain within consensus blocks. Validators are tightly co-located, which keeps consensus latency low enough to support order-book trading workflows familiar from centralized exchanges. Trading is fee-funded rather than gas-funded, eliminating one source of execution variance present on other on-chain venues.

The protocol's matching rules differ from a standard continuous limit order book in one substantive way. Within each block, the validator software processes cancel orders and post-only limit orders before processing good-til-cancel and immediate-or-cancel orders. When the underlying spot price moves, a market maker with resting quotes faces an immediate adverse-selection threat, since any taker arriving in the same block is more likely to be informed about the price move than the average trader. The maker's natural response is to cancel and re-quote. Without cancel prioritization, a sufficiently fast taker can submit a marketable order in the same block as the maker's cancel, and have it processed first - the maker is picked off at a stale price. Cancel prioritization eliminates this specific failure mode.

The Hyperliquid Foundation publishes a daily archive of L2 order-book snapshots to a public S3 bucket (Hyperliquid Foundation, 2024). As measured in our sample, snapshots are produced at a median inter-snapshot gap of 0.543 seconds, yielding approximately 110 snapshots per asset per minute. Each snapshot contains the top 20 levels on each side of the book. Trade data are available separately via the Hyperliquid API but are not included in the public S3 archive; the analysis in this paper uses only book-snapshot data and does not require trade-level inputs.

### 3.2. The October 10 Cascade and the November 3-4 Selloff

On October 10, 2025, the cryptocurrency perpetual futures market experienced a system-wide liquidation cascade. Across centralized and decentralized venues, leveraged positions were forcibly closed at scale; aggregate liquidation volume across all major venues exceeded nineteen billion US dollars within a 24-hour window. Hyperliquid's automated deleveraging mechanism was triggered for several contracts, and the venue's inside market widened sharply for a multi-hour window beginning approximately 21:00 UTC.

Twenty-five days after the cascade, on November 3-4, 2025, the market experienced a smaller but still substantial selloff. Bitcoin declined from approximately $107,000 on November 1 to $99,000-$104,000 on November 4, with daily moves in excess of five percent and renewed liquidations across major perpetual venues. The November event is an order of magnitude smaller than October 10 in absolute liquidation volume but is large enough to produce visible widening in our post-event recovery trajectory. We treat the cascade as the primary event and the November selloff as a confound to be explicitly controlled for in the empirical analysis.

### 3.3. Data and Sample Construction

Our sample covers September 15 through December 15, 2025, a 92-day window centered on the cascade event. We restrict the asset universe to the five perpetual contracts with continuous existence on Hyperliquid throughout the sample and with sufficient daily volume to support meaningful inside-market inference, namely BTC, ETH, SOL, XRP, and DOGE.

We obtain Hyperliquid L2 order-book snapshots from the public archive bucket. The bucket is requester-pays, meaning the data are free at source but the requester bears the egress cost; total data-transfer cost for our sample was approximately ten US dollars. Snapshots are stored as line-delimited JSON within hourly LZ4-compressed files. We extracted, parsed, and re-stored each asset-day as a Parquet file, retaining the original timestamps and the full twenty-level book on each side.

From each L2 snapshot we compute the quoted spread at the inside (in basis points relative to the midpoint), the inside-quote depth in US dollars (the dollar size resting at the best bid plus the dollar size resting at the best ask), and the simulated round-trip execution cost for marketable orders at $50K, $100K, and $500K notional sizes. Realized volatility is computed as the annualized standard deviation of one-second log midpoint returns within each minute, expressed in basis points. All metrics are aggregated to one-minute resolution by taking the first value within each minute bin.

After construction, the venue-internal panel contains 531,305-minute observations across 78 dates and 5 assets. Coverage is essentially complete in the pre-event baseline window (October 3-9, 2025: 10,080 minutes per asset) and in the primary post-event window (October 11 through November 2, 2025: approximately 25,000 minutes per asset accounting for a six-day archive gap October 17-22). The Hyperliquid public archive also published no L2 snapshots for November 18-25, 2025; observations for those date ranges are missing from the extended recovery analysis.

### 3.4. Methodology

The analysis proceeds in five steps. First, we characterize the venue's normal-state inside market using the pre-event baseline window. Second, we compute hour-by-hour inside-market statistics on October 10 UTC to characterize the intraday cascade evolution. Third, we compute daily median spread and depth ratios relative to the pre-event baseline for each asset across the 51-day post-event window, smoothed with seven-day rolling medians. Fourth, we re-run the recovery analysis restricted to the 23-day clean window (October 11 to November 2, 2025) that excludes the November 3-4 selloff, to verify that the spread and depth findings are attributable to the cascade rather than to the subsequent event. Fifth, we compute time-to-threshold recovery measures for spreads at 1.5x, 1.2x, and 1.1x of pre-event baseline.

We define the metrics used in the paper as follows. For a given L2 snapshot, let $a$ and $b$ denote the best (highest) bid price and best (lowest) ask price, with associated resting sizes $q_b$ and $q_a$ in base units. The midpoint is $m = (a + b)/2$. The quoted spread in basis points is $S = 10{,}000 \times (a - b)/m$. The inside-quote depth in US dollars is $D = (q_b \times b) + (q_a \times a)$, the dollar size resting at the best bid plus that resting at the best ask. Snapshot-level values are aggregated to one-minute resolution by taking the first valid observation in each minute bin, yielding minute series $S_t$ and $D_t$ for each asset.

For each asset, the pre-event baseline median spread is the median of $S_t$ over all minute bins in the baseline window (October 3-9, 2025), denoted $\overline{S}_0$; the baseline median depth $\overline{D}_0$ is defined analogously. The daily median spread on day $d$ is the median of $S_t$ over the minute bins in that day, $M_S(d) = \text{median}\{S_t : t \in d\}$, and the daily median depth $M_D(d) = \text{median}\{D_t : t \in d\}$. The reported spread ratio and depth ratio on day $d$ are $R_S(d) = M_S(d)/\overline{S}_0$ and $R_D(d) = M_D(d)/\overline{D}_0$, so that a value of 1 corresponds to the pre-event baseline. Reported series are smoothed with a centred seven-day rolling median of the daily ratios. Percentiles (for example the 95th percentile of spread, $p_{95}$) are computed as the empirical percentile of the minute-level $S_t$ within the relevant hour or day bin, and the maximum is the largest minute-level $S_t$ in that bin. Realized volatility within a minute is the annualized standard deviation of one-second log midpoint returns in that minute, expressed in basis points. All analysis code and derived data products are made publicly available; see the Data Availability Statement.

## 4. Empirical Results and Discussion

### 4.1. Pre-Event Inside-Market Characterization

The cross-sectional pattern is consistent with stylized facts about market quality across asset capitalizations. BTC quotes a tight inside spread (median 0.088 bps) backed by approximately $1.8 million of inside depth; DOGE quotes a wider spread (0.562 bps) backed by approximately $28,000 of inside depth. The five assets span roughly an order of magnitude in pre-event spread and approximately three orders of magnitude in pre-event depth. On a 24/7 venue with no opening or closing auctions, median quoted spreads pooled across assets vary by only seven percent across UTC hours of the day and six percent across days of the week, contrasting sharply with conventional equity venues (Wood et al., 1985). The flatness is itself an institutional feature of continuous global trading without cohort handovers. Table 1 reports the pre-event baselines across the five assets.

**Table 1.** Pre-event inside-market baselines (October 3-9, 2025).

| Asset | Median spread (bps) | Mean spread (bps) | Median depth (USD) | Median realized vol (bps) |
|---|---|---|---|---|
| BTC | 0.088 | 0.097 | 1,825,661 | 2.11 |
| ETH | 0.238 | 0.248 | 1,400,559 | 3.65 |
| SOL | 0.467 | 0.479 | 444,653 | 5.57 |
| XRP | 0.369 | 0.385 | 77,696 | 4.49 |
| DOGE | 0.562 | 0.602 | 28,135 | 6.46 |

Note: Each cell is computed across 10,080-minute observations per asset (7 days × 1,440 minutes). Inside depth is the dollar value resting at the best bid plus the best ask, summed across both sides. Realized volatility is the annualized standard deviation of one-second log midpoint returns within each minute, in basis points.

### 4.2 Intraday Evolution on October 10

Figure 1 presents the minute-level evolution of inside spread on October 10, 2025 UTC for each of the five assets. Spreads on all five assets remained near baseline through approximately hour 13. A first elevation is visible around hour 15 (corresponding to a brief BTC drawdown), partial recovery through hour 17, and a second elevation around hour 19. The principal cascade peak occurred in hour 21 UTC, with all five assets simultaneously experiencing spread widening of one to two orders of magnitude.

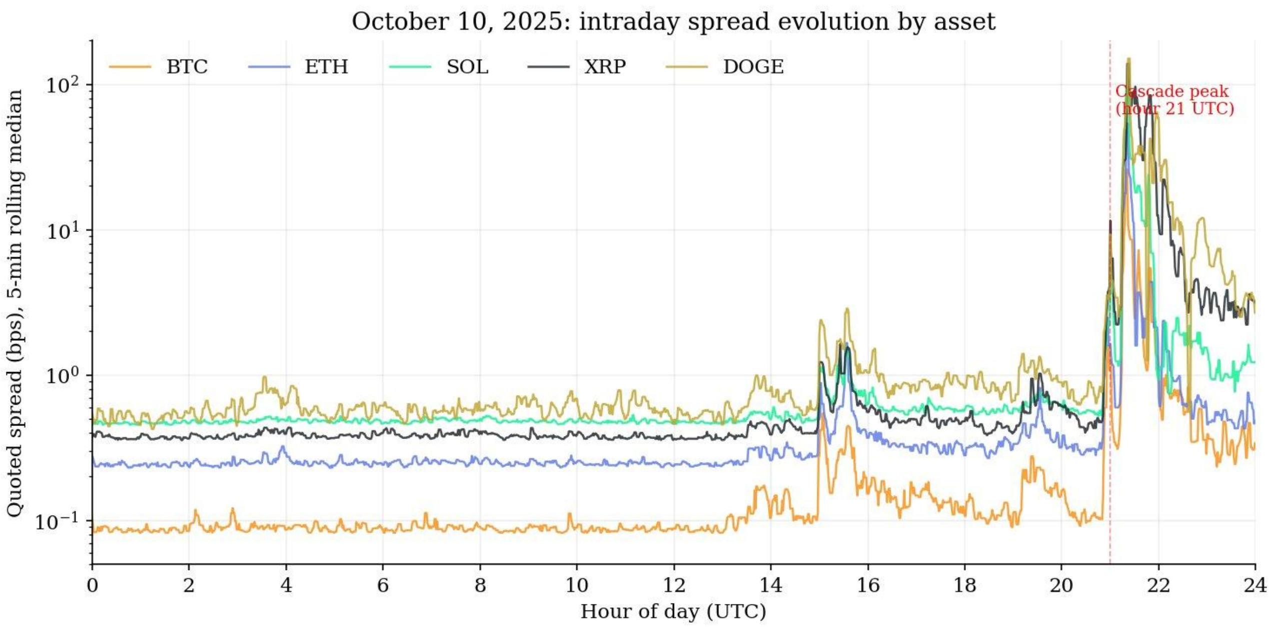


**Figure 1.** Minute-level inside spread on October 10, 2025 UTC, by asset. Five-minute rolling median, log scale. Vertical dashed line indicates the principal cascade peak at hour 21 UTC.

Table 2 reports the hour-by-hour evolution pooled across assets.

**Table 2.** Hour-by-hour inside-market evolution, October 10, 2025 UTC.

| Hour UTC | Median spread (bps) | p95 spread (bps) | Max spread (bps) | Median depth (USD) | Median rv (bps) |
|---|---|---|---|---|---|
| 00 | 0.372 | 0.566 | 0.639 | 415,716 | 3.43 |
| 04 | 0.378 | 0.667 | 1.014 | 604,670 | 4.59 |
| 08 | 0.381 | 0.658 | 0.904 | 520,120 | 4.37 |
| 12 | 0.368 | 0.577 | 0.728 | 493,769 | 3.44 |
| 13 | 0.399 | 0.802 | 1.071 | 548,614 | 6.64 |
| 14 | 0.441 | 0.802 | 4.135 | 597,182 | 7.27 |
| 15 | 0.724 | 1.983 | 4.968 | 463,643 | 16.60 |
| 16 | 0.497 | 1.009 | 1.568 | 623,944 | 9.89 |
| 17 | 0.485 | 0.936 | 1.188 | 487,329 | 8.99 |
| 18 | 0.439 | 0.960 | 1.178 | 515,262 | 8.04 |
| 19 | 0.588 | 1.291 | 2.852 | 566,483 | 12.78 |
| 20 | 0.534 | 3.030 | 13.692 | 514,913 | 8.66 |
| 21 | 9.376 | 87.476 | 241.819 | 391,473 | 150.62 |
| 22 | 1.709 | 18.804 | 44.351 | 375,291 | 42.46 |
| 23 | 1.080 | 5.576 | 8.713 | 217,889 | 31.14 |

Note: Statistics are pooled across all five assets within each UTC hour bin. Hours 1, 2, 3, 5, 6, 7, 9, 10, 11 are similar to hour 0 and omitted for brevity.

The first 14 hours of the day were essentially indistinguishable from a normal trading day. Beginning hour 15, both spread and realized volatility elevated noticeably, and through hours 16 to 20 the inside market was moderately stressed. The peak occurred in hour 21, when median spread reached 9.38 bps (more than 25 times normal levels), the 95th percentile reached 87.5 bps, the maximum minute observation reached 241.8 bps, and median realized volatility reached 150.6 bps (more than 30 times normal). By hour 22 the median spread had already declined to 1.71 bps, and by hour 23 to 1.08 bps.

### 4.3 Spread Recovery in the Clean Post-Event Window

Because a separate market-wide selloff began on November 3, 2025, our primary recovery analysis focuses on the 23-day clean window between the cascade and the second event, that is, October 11 through November 2, 2025. Table 3 reports the key recovery metrics in this window.

**Table 3.** Recovery metrics in the clean post-event window (October 11 to November 2, 2025).

| **Asset** | **Max spread ratio** | **End-of-window spread ratio** | **Days to ≤1.2x** | **Days to ≤1.1x** |
|---|---|---|---|---|
| BTC | 1.31x | 1.07x | 6 | 15 |
| ETH | 1.41x | 1.13x | 6 | 15 |
| SOL | 1.43x | 1.19x | 14 | Not reached |
| XRP | 1.82x | 1.14x | 15 | Not reached |
| DOGE | 1.77x | 1.06x | 13 | 14 |

Note: Spread ratio is the daily median quoted spread relative to the pre-event baseline (October 3-9, 2025). "Max" refers to the maximum daily median observed during the 23-day window. "End-of-window" is the spread ratio on November 2, 2025. Threshold-recovery columns report the number of days after October 10 required to first reach the threshold; "Not reached" indicates the threshold was not crossed within the 23-day window.

The recovery in quoted spread is rapid and close to complete. Within 6 days, BTC and ETH were within 1.2 times pre-event spread; within 13-15 days, BTC, ETH, and DOGE were within 1.1 times. By the end of the 23-day window, BTC (1.07x) and DOGE (1.06x) were essentially fully recovered, ETH (1.13x) was close, and the two laggards (SOL at 1.19x and XRP at 1.14x) were within approximately 20 percent. This pattern contrasts with the more common stylized finding that stress-event spread widening persists for extended periods; on Hyperliquid, the spread dimension of execution quality was substantially restored within a month of a major cascade.

Figure 2 presents the seven-day rolling spread ratio for each asset across the full 51-day post-event window.

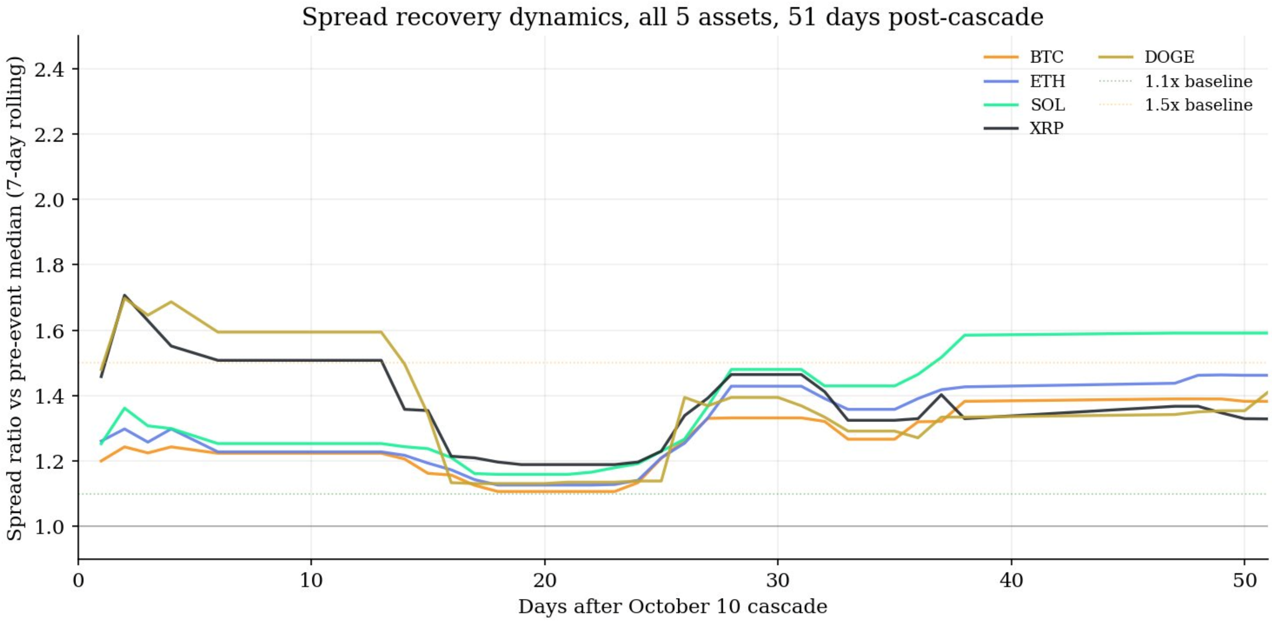


**Figure 2.** Daily spread ratio relative to pre-event baseline, seven-day rolling median, by asset. Days 1-23 correspond to the clean post-event window. Days 24 onward are confounded by the November 3-4 market-wide selloff.

### 4.4 Depth Contraction - The Persistent Dimension

The recovery story in depth is qualitatively different from the spread story. Table 4 reports depth dynamics in the same clean post-event window.

Four of five assets (all except XRP) showed further depth contraction from the first week to the third week of the clean post-event window. BTC depth declined by 16.2 percentage points, DOGE by 14.7 percentage points, ETH by 9.8 percentage points, SOL by 2.7 percentage points. Only XRP showed any depth recovery, and the improvement was marginal (+1.5 pp) from an already-low starting level (0.76x). Four of five assets ended the clean window with depth materially below pre-event levels, and three were below 0.80x on the days 17-23 median measure. The depth contraction preceded the November 3-4 event and is therefore attributable to the cascade itself, not to the subsequent disturbance.

**Table 4.** Depth contraction in the clean post-event window (October 11 to November 2, 2025).

| Asset | Days 1-7 median depth ratio | Days 17-23 median depth ratio | Change (percentage points) |
|---|---|---|---|
| BTC | 0.93x | 0.77x | -16.2 |
| ETH | 0.88x | 0.79x | -9.8 |
| SOL | 0.93x | 0.90x | -2.7 |
| XRP | 0.76x | 0.78x | +1.5 |
| DOGE | 0.97x | 0.83x | -14.7 |

Note: Depth ratio is the daily median inside-quote depth relative to the pre-event baseline. “Days 1-7 median” aggregates October 11-16 (with October 17 missing due to archive gap). “Days 17-23 median” aggregates October 27 to November 2.

Figure 3 extends the depth analysis through the full 51-day post-event window. The pattern differs markedly from the spread figure, in that depth trajectories are monotonically declining through the entire window with no visible second-event signature. Two contracts (ETH and DOGE) cross the 0.5x line by the end of the observation period — inside depth cut by half relative to pre-event seven weeks later.

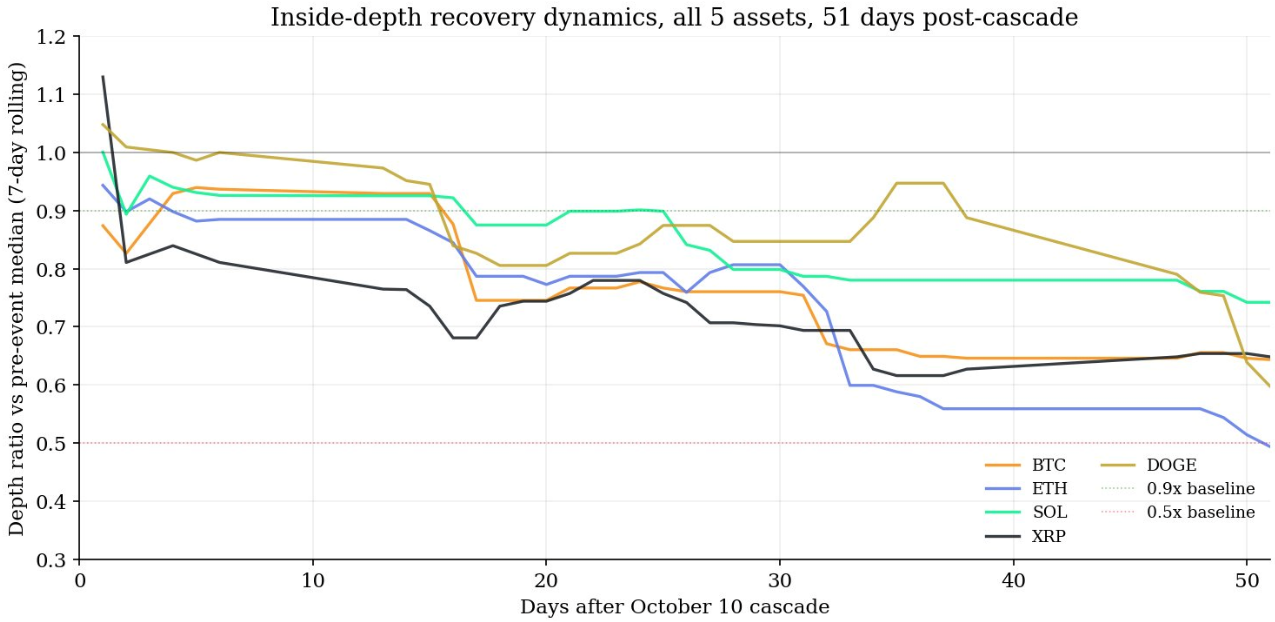


**Figure 3.** Daily depth ratio relative to pre-event baseline, seven-day rolling median, by asset. Unlike the spread series in Figure 2, depth shows monotonic deterioration over the entire 51-day window, with no visible second-event inflection.

### 4.5 Cross-Sectional Pattern

Figure 4 presents the cross-sectional pattern at 30 days post-cascade. The left panel plots spread widening against pre-event spread; the right panel plots depth contraction against the same.

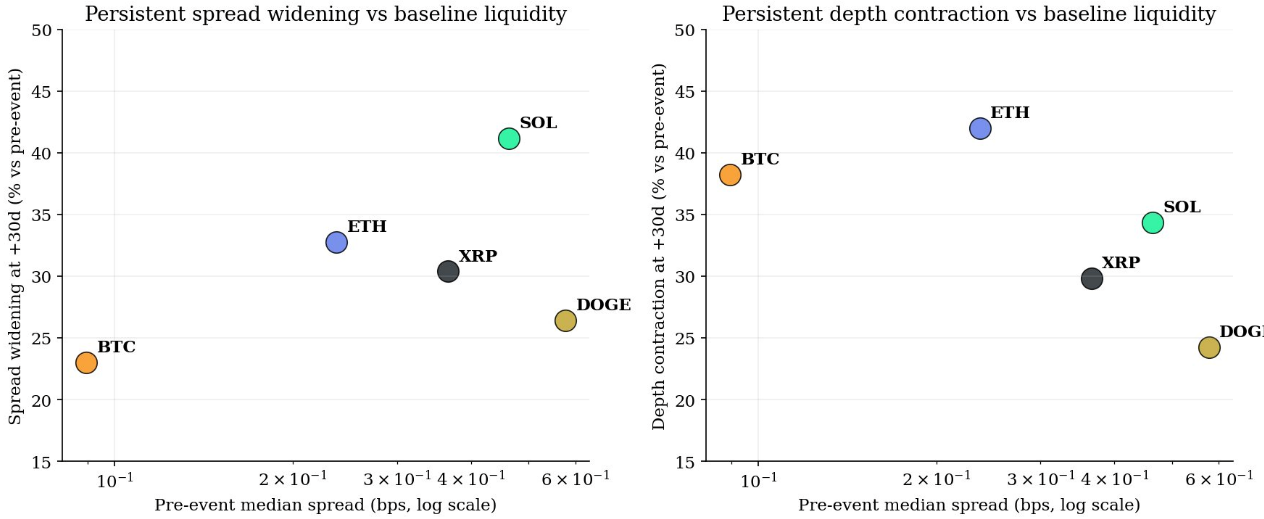


**Figure 4.** Cross-sectional pattern of persistent post-event effects at 30 days. Left: spread widening (% vs pre-event). Right: depth contraction (% vs pre-event).

The depth cross-section is the cleaner finding and is interpretable as follows. BTC and ETH (the two largest contracts by pre-event depth) exhibit the largest depth contractions (38 and 42 percent). The pattern is consistent with an interpretation in which tail-risk perception increases proportionally to notional exposure, in that makers committed the most inventory to the largest contracts pre-event and therefore had the most inventory to reduce post-event. The smallest contract (DOGE) exhibits the smallest percentage depth contraction, but this should not be interpreted as DOGE being least affected — DOGE had the least depth to begin with.

### 4.6 Discussion: Two-Regime Recovery

The central empirical finding of this paper is the two-regime recovery pattern, in which pricing recovers quickly and nearly completely within three weeks of a cascade while inventory commitment contracts persistently over the same window and continues to deteriorate over the full observation period. The pattern is consistent with the inventory-management framework of Stoll (1978) and Ho and Stoll (1981), in which market makers set both quoted prices and committed sizes as separable decisions responding to different risk considerations. A tail event raises the perceived variance of inventory holdings, and the rational maker response is to reduce size exposure rather than demand a wider spread - particularly in a competitive market where quoted prices are disciplined by the possibility of being undercut.

The pattern has implications for how venue resilience should be measured. A metric that examines only the spread dimension - as is conventional in most empirical microstructure work - will systematically understate the post-event impairment of venue liquidity when makers respond primarily through size reduction. A complete resilience assessment requires monitoring both dimensions, with particular attention to depth which the evidence here suggests is the more persistent and potentially more consequential dimension for large-order execution quality.

Hyperliquid's distinguishing market-design feature is the consensus-layer cancel prioritization rule, motivated by maker protection. The findings are consistent with this protection operating during normal-state trading, in that pre-event spreads are tight relative to the venue's age and intraday patterns are flat (Section

4.1). The rapid spread recovery documented in Section 4.3 is also consistent with an environment where makers feel confident re-entering quoted prices quickly post-stress. But the persistent depth contraction (Section 4.4) suggests that cancel prioritization, while beneficial for normal-state market quality and for fast pricing recovery, does not insulate the venue from the inventory-commitment effects of a tail event. Venue resilience and venue-design quality are distinct properties.

There are some limitations. First, the Hyperliquid public archive publishes only the top twenty levels of the book on each side. Wider-book dynamics (depth at, say, 25 basis points outside the inside) are not directly observable from the public data; the depth metric we report captures top-of-book commitment specifically. Second, the archive published no L2 snapshots for October 17-22 and November 18-25, 2025. The first gap falls within the clean post-event window and reduces our coverage of the 7-12 day post-cascade range. Third, we cannot directly identify individual market makers; our analysis is silent on whether the persistent inventory contraction reflects a few large makers exiting or many small makers each reducing exposure marginally. Fourth, the November 3-4 selloff confounds post-event analysis beyond day 23; our primary recovery findings are restricted to the clean window. Fifth, the analysis is venue-internal by design and does not address whether Hyperliquid's recovery dynamics are typical or atypical relative to centralized perpetual venues experiencing the same cascade.

## 5. Conclusion

This paper documents the inside-market response of Hyperliquid - the largest decentralized perpetual futures venue - to the October 10, 2025 system-wide liquidation cascade. Using the venue's public L2 archive, we construct a minute-level panel spanning five major perpetual contracts across a 92-day window centered on the event. The analysis yields three main findings.

First, the cascade was a tightly compressed intraday event. Median quoted spread pooled across assets reached 9.4 basis points at hour 21 UTC and the 95th percentile reached 87.5 basis points, with intraday recovery well under way by end of day. Second, in the 23-day clean post-event window prior to a separate market-wide selloff, quoted spreads recovered substantially - three of five contracts within 1.1 times pre-event levels and all within 1.2 times. Third, inside-quote depth contracted persistently over the same clean window, with four of five contracts showing further depth reduction from days 1-7 to days 17-23 post-event and none fully recovering over the full 51-day observation window.

The contrast between the two findings - spread recovery near complete while depth contraction persists - motivates what we term two-regime liquidity recovery. Market makers maintain competitive quoted prices but reduce the inventory they commit at those quotes following a tail event. Three implications follow. For microstructure theory, the finding is consistent with the classic inventory-management framework's prediction that pricing and sizing are separable maker decisions responding to different risk considerations. For empirical market-resilience research, the finding implies that spread-only metrics systematically understate the impairment of venue liquidity following stress events and that depth metrics should be a standard component of resilience analysis.

For decentralized derivatives venue design, the finding establishes that a venue specifically engineered to protect quoting market makers can still experience multi-week impairment of inventory commitment after a tail event; venue resilience and venue-design quality are distinct properties. For investors and traders, the finding carries a practical implication. In the weeks after a cascade, a tight quoted spread can give a misleading impression that liquidity has normalized, while the depth available at the inside remains materially reduced. Investors who size orders against the quoted spread alone risk underestimating the price impact and execution cost of larger trades during this recovery phase. Those executing size in the post-event window may benefit from monitoring inside depth directly, working orders more patiently, or scaling trade sizes to the reduced depth rather than to the apparently recovered spread.

Beyond the specific findings, this paper also takes a methodological position. The public order-book archives that on-chain venues publish offer opportunities for transparent and replicable empirical microstructure research that are difficult to match in centralized venues that typically restrict tick-level access. We hope the results encourage future replication and extension on other on-chain venues with similar archive practices.

**Data Availability Statement**

All code, derived data products, and figures used in this study are publicly available at the project repository: https://github.com/boonchuan/hyperliquid. The raw Hyperliquid L2 order-book snapshots analyzed are publicly available from Hyperliquid's S3 archive at s3://hyperliquid-archive (requester-pays bucket; total egress cost for the full sample was approximately ten US dollars). The repository contains the full analysis pipeline - data-pull scripts, metrics construction, cascade event study, and figure generation - that reproduces every table, figure, and numerical claim in this paper from the public raw data.

**Funding**

This research received no external funding. The author self-funded all data-access costs (approximately ten US dollars in AWS egress fees) and computational resources used in this study.

**Acknowledgements**

The author thanks the Hyperliquid Foundation for maintaining the public L2 order-book archive that made this study possible.

**Declaration of the Use of Generative AI**

During the preparation of this work, the author used Claude (Anthropic) for drafting prose, editing, generating Python code for data analysis, and structuring the manuscript. All empirical analysis, data pulls, methodology choices, and interpretations are the author's own. Every number, table, and figure in the paper was generated from the author's own data pipeline running on the author's machine; AI-generated text was verified against the underlying data and revised by the author throughout. The AI tool was not used to generate citations or references; all references were verified by the author against primary sources. The AI system is not an author and bears no responsibility for the content. After using this tool, the author reviewed, revised, and verified all content to ensure accuracy and accepts full responsibility for all aspects of the publication.

## Conflicts of Interest

The author declares no conflicts of interest. The author has no financial relationship with, position at, or advisory role for Hyperliquid, the Hyperliquid Foundation, Binance, or any of the venue operators referenced in this paper. The author is an independent researcher and received no funding from any party with an interest in the findings.

## References

Anand, A., & Venkataraman, K. (2016). Market conditions, fragility, and the economics of market making. *Journal of Financial Economics*, 121(2), 327–349. https://doi.org/10.1016/j.jfineco.2016.03.006

Bessembinder, H., Hao, J., & Zheng, K. (2015). Market making contracts, firm value, and the IPO decision. *Journal of Finance*, 70(5), 1953–2002. https://doi.org/10.1111/jofi.12285

Brauneis, A., Mestel, R., Riordan, R., & Theissen, E. (2021). How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance*, 124, 106041. https://doi.org/10.1016/j.jbankfin.2020.106041

Brogaard, J., Carrion, A., Moyaert, T., Riordan, R., Shkilko, A., & Sokolov, K. (2018). High frequency trading and extreme price movements. *Journal of Financial Economics*, 128(2), 253–265. https://doi.org/10.1016/j.jfineco.2018.02.002

Chitra, T. (2025). Autodeleveraging: Impossibilities and optimization [Working paper].

Foley, S., Karlsen, J. R., & Putnins, T. J. (2019). Sex, drugs, and Bitcoin: How much illegal activity is financed through cryptocurrencies? *Review of Financial Studies*, 32(5), 1798–1853. https://doi.org/10.1093/rfs/hhz015

Foucault, T., Kadan, O., & Kandel, E. (2013). Liquidity cycles and make/take fees in electronic markets. *Journal of Finance*, 68(1), 299–341. https://doi.org/10.1111/j.1540-6261.2012.01801.x

Hasbrouck, J., & Sofianos, G. (1993). The trades of market makers: An empirical analysis of NYSE specialists. *Journal of Finance*, 48(5), 1565–1593. https://doi.org/10.1111/j.1540-6261.1993.tb05121.x

He, S., Manela, A., Ross, O., & von Wachter, V. (2024). Fundamentals of perpetual futures [Working paper]. Washington University in St. Louis.

Hendershott, T., & Moulton, P. C. (2011). Automation, speed, and stock market quality: The NYSE's Hybrid. *Journal of Financial Markets*, 14(4), 568–604. https://doi.org/10.1016/j.finmar.2011.02.003

Ho, T., & Stoll, H. R. (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1), 47–73. https://doi.org/10.1016/0304-405X(81)90020-9

Hyperliquid Foundation. (2024). Latency and transaction ordering on Hyperliquid. https://hyperliquid.gitbook.io

Madhavan, A., & Smidt, S. (1991). A Bayesian model of intraday specialist pricing. *Journal of Financial Economics*, 30(1), 99–134. https://doi.org/10.1016/0304-405X(91)90023-D

Ruan, Q., & Streltsov, A. (2024). The emerging market of cryptocurrencies and perpetual contracts [Working paper]. Cornell University.

Stoll, H. R. (1978). The supply of dealer services in securities markets. *Journal of Finance*, 33(4), 1133–1151. https://doi.org/10.1111/j.1540-6261.1978.tb02053.x

Wood, R. A., McInish, T. H., & Ord, J. K. (1985). An investigation of transactions data for NYSE stocks. *Journal of Finance*, 40(3), 723–739. https://doi.org/10.1111/j.1540-6261.1985.tb04996.x